1988

KNOWLEDGE ACQUISITION METHODOLOGIES: SURVEY AND EMPIRICAL ASSESSMENT

Martha Grabowski
Rensselaer Polytechnic Institute

KNOWLEDGE ACQUISITION METHODOLOGIES: SURVEY AND EMPIRICAL ASSESSMENT

Martha Grabowski
Decision Sciences and Engineering Systems
Rensselaer Polytechnic Institute

ABSTRACT

Knowledge acquisition, the process of extracting information from human experts, is one of the challenges in building expert systems. Modern practitioners and researchers need more guidance than is provided by existing knowledge acquisition guidelines. However, there has been little empirical research upon which to base the needed guidelines. This paper surveys the available knowledge acquisition techniques and describes a knowledge acquisition experiment which contrasts three of these methods. A framework was developed to categorize the types of heuristic which can be elicited with different means of knowledge acquisition. This research represents the initial steps in a research program focused on the development of empirically evaluated, generalized guidelines for effecting knowledge acquisition.

1. INTRODUCTION

Several tasks are important in building expert systems. The first is that of knowledge acquisition, the process by which an expert's problem-solving knowledge is identified, elicited, modeled, and incorporated into an expert system. Current methods for transferring human expertise into a knowledge base are time-consuming, expensive, and constitute a bottleneck in the expert system development process (Buchanan et al. 1983). Butler and Corter (1986) suggest that the problem with knowledge acquisition is due in large part to two factors: heavy reliance by the knowledge engineer on unstructured interviews with the experts and heavy reliance on introspection by the expert. The interviews require a great deal of time and patience from people with rare skills. The development of effective interviewing skills for knowledge engineers (and traditional systems analysts) requires substantial training and practice. Experts may not be able to introspect accurately and the validity of introspective reports is the topic of much discussion. Using current methods, building a reasonably-sized expert system requires several man-years of effort (Butler and Corter 1986; Waterman 1986).

Thus, there is a need for a knowledge acquisition methodology which spans both research and commercial environments, and which is amenable to a variety of implementations. This research examines several knowledge acquisition methodologies, provides some empirical support for the use a combination of knowledge acquisition techniques, and lays the foundation for a pragmatic knowledge acquisition methodology, long needed by information systems and expert systems researchers and practitioners.

2. KNOWLEDGE ACQUISITION METHODS

Today's knowledge engineers may use a pad and pencil, a tape recorder, or even a videotape to record an expert's thoughts, ideas, concepts, or responses to questions. One approach to knowledge acquisition is known as "analyst becomes expert" (Brooking 1986). The ACE system, developed at Bell Laboratories (Vesonder et al. 1983) adopted this approach. An analyst who had experience in a technical area (in this case, a psychologist who in the past worked as a telephone engineer) was trained to acquire expertise in the chosen domain.

Initial prototypes were built based on the knowledge acquired and were refined using a "real" expert as a knowledge source. The reverse approach to knowledge acquisition is the "expert becomes analyst" approach (Brooking 1986). Probably the most famous system built under this approach is the MYCIN system built by Buchanan and Shortliffe (1984).

In the 1950s and 1960s, the main emphasis of most artificial intelligence programs was demonstrating intelligent behavior for a few limited problems. Programmers acted as their own experts and coded in the domain expertise. Today, for most problems, the programmers and the expert are not the same person and it is risky to rely on such "handcrafting" to build complex programs embodying large amounts of judgmental information (Buchanan and Shortliffe 1984).

2.1 Interviews

Reviews of knowledge acquisition techniques (Gammack and Young 1985; Raulefs 1985) indicate that the use of written materials describing a particular domain of expertise, coupled with verbal interviews with experts, serve as the dominant means of gathering expert judgment. A typical scenario starts with a dialogue between the knowledge engineer and the expert. The relevant concepts in the problem domain are identified and the relationships between these concepts are made explicit (Hayes-Roth, Waterman, and Lenat 1983; Buchanan and Shortliffe 1984). After initial conceptualization, in which most of the framework for talking about the subject matter is designed, the knowledge structure can be filled in rather rapidly. This effort is then generally followed by testing and refinement of the knowledge base.

Interviews may take two forms: structured and unstructured (Fellers 1987). In an *unstructured interview*, the expert often performs a familiar task, one that he/she performs on a frequent basis, while the knowledge engineer asks "more or less spontaneous questions" (Hoffman 1987). These extensive interviews often last months and may even take place over several years. During the unstructured interview, the knowledge engineer actively questions the expert, who is consciously focusing on the knowledge being used in the problem-solving process. Many knowledge engineers rely solely on the unstructured interview (Fellers 1987). *Structured interviews*, in contrast, take place after the initial knowledge base has been established, and are used to refine the knowledge base. Fellers suggests that limited information tasks may be used during structured interviews to restrict the amount of information elicited and to gain additional knowledge about how an expert performs a task.

2.2 Constrained Processing Tasks

In addition to interviewing the expert, knowledge engineers often ask experts to perform constrained processing tasks (Fellers 1987). Hoffman (1987) discusses two methods of forcing the expert to focus more on actual problem-solving processes: simulated familiar tasks and scenarios. *Simulated familiar tasks* allow an expert to simulate performance of a familiar task using a variety of tools--a case study (Buchanan et al. 1983; Rolandi 1986), a simulation (Prerau 1987; Grabowski 1987), "tough cases" (Hoffman 1987), or an expert system prototype (Grabowski 1987). Prerau discusses using simulated familiar tasks for a system he developed, using both hand and computer simulation. A variation on this approach is *actual familiar tasks*, where an expert performs actual physical tasks within the problem domain which the knowledge engineer observes (Grabowski 1987).

When using *scenarios*, experts imagine typical scenarios set for them in the problem domain by the knowledge engineer. The expert draws on analogies of previous situations. Scenarios can be thought of as "what if" kinds of analyses, in that they force the expert to concentrate on specific aspects of a case or problem (Fellers 1987). Bimson and Burris (1987) found scenarios to be a critical component of their knowledge-gathering activity. Fellers provides a comprehensive review of different knowledge acquisition techniques and advocates the use of a primary and a secondary knowledge acquisition method.

2.3 Protocol Analysis

Protocol analysis is the process of translating verbalizations of "thinking aloud" subjects into more accessible and meaningful representations (Bouwman 1978). Subjects are asked to "think aloud" while solving a problem or making a decision and the translation of these verbalizations provides the researcher with a sequenced map of the expert's decision-making thoughts (Newell and Simon 1972; Ericsson and Simon 1980; Schweiger 1983). The verbalizations are tape-recorded, transcribed, and content analyzed via a coding scheme. *Concurrent protocols* are obtained by recording the expert's thinking aloud thoughts at the same time the expert solves the problem, while *retrospective protocols* are obtained by asking the expert to review records (audio, video, transcriptions) of the expert's verbalizations after the task is completed. Retrospective protocols are often used when concurrent protocols are suspected to affect the expert's task performance or when task performance is suspected to interfere with the expert's ability to offer a coherent protocol (Wright and Ayton 1987).

Context-focusing (Wright and Ayton 1987), or short-cut protocol analysis, is another technique that gives the knowledge engineer access to the expert's sequence of rule testing. In context-focusing, the knowledge engineer imagines a particular domain state and the expert has to find out what it is by querying, much in the manner of a "twenty questions" game. The knowledge engineer initiates the procedure several times, each time imagining alternative system states, which allows the knowledge engineer to examine the expert's priority ordering of rules and objects in the domain.

2.4 Declarative Knowledge Elicitation

A number of knowledge acquisition techniques have been proposed which incorporate techniques from other disciplines: multidimensional scaling (Eliot 1986; Whalley 1984), network scaling (Whalley 1984), cluster analysis (Cooke and McDonald 1986), discourse analysis (Belkin, Brooks, and Daniels 1986), and psychological scaling (Cooke and McDonald 1986). With *multidimensional scaling* (MDS), experts are asked to rate the similarity of objects and represent the similarities as distances on a seven-point scale ranging from no similarity to completely similar. The intention is to determine the expert's rank ordering of objects within a problem domain. *Card sorting techniques* (Wright and Ayton 1987), short-cuts to eliciting declarative knowledge, involve having the knowledge engineer write the names of objects, experiences, and rules in the expert's world onto cards. Only those concepts which the knowledge engineer feels need to be explored are used in the card sorting tasks. The expert then sorts the cards, and the expert's card ordering provides a map of the expert's domain classifications and relationships. The cards also provide a mechanism for understanding the expert's domain jargon and its underlying structure.

2.5 Automated Knowledge Acquisition

Recent experiences in knowledge acquisition have used knowledge acquisition software or automated knowledge engineering tools. This software may allow users to add new knowledge to the knowledge base and then check the resulting knowledge base for consistency or reasonableness (Davis, Buchanan, and Shortliffe 1977), or allow the knowledge base to be built with domain information and examples of the expert's decisions, using the system to determine the knowledge base's general rules (Greene 1987). Knowledge engineering software may also provide explanations of how the system's conclusions were arrived at.

Automated knowledge acquisition tools have yet to withstand the tests of time and rigorous empirical research (Greene 1987). However, use of automated knowledge acquisition techniques lets the knowledge engineer combine the traditional advantages of rapid prototyping with the additional advantages of streamlined knowledge acquisition (Grabowski 1987).

2.6 Prototyping for Knowledge Acquisition

The evolutionary approach of constructing and testing increasingly more elaborate prototypes is the most prevalent strategy for designing expert systems (Waterman 1986; Davis et al. 1981). Hayes-Roth, Waterman, and Lenat suggest that interactive prototyping is also an effective knowledge acquisition technique. With prototyping, the knowledge engineer and the expert work together to quickly identify the basic information requirements and to build a prototype expert system quickly. The prototype is then used to refine the knowledge base requirements, and the prototype is revised and enhanced, iteratively, as necessary (Janson and Smith 1985; Davis and Olson 1985).

Prototyping facilitates the knowledge acquisition process by helping the knowledge engineer extract rules from the expert (Grabowski and Wallace 1986). Once an initial prototype is constructed, both the knowledge engineer and the expert have a common base of reference from which to develop the remainder of the system. Since the prototype uses the same terminology that the expert uses to solve problems in the domain, the knowledge engineer should have an easier time correcting and modifying the system's rules and relationships (Alavi 1984). Use of the prototype also counters waning expert enthusiasm, which often develops in later stages of expert system development work.

2.7 Knowledge Acquisition to Knowledge Representation

Once the expert's information is elicited, there are two methods for molding the knowledge into a form appropriate for knowledge base representation: *rapid prototyping*, which mixes the knowledge acquisition and implementation stages, and *structured knowledge acquisition* (deGreef and Breuker 1985), which separates the knowledge acquisition and implementation stages. This second methodology is not widely practiced. In the first method, the knowledge engineer uses interview data from human experts and immediately starts to build a prototype in an implementation formalism.

In structured knowledge acquisition, knowledge acquisition and implementation are separated. The task for the knowledge engineer is to bridge the gap between the verbal data from the experts and the actual implementation of the system. Crucial to this methodology is the use of thinking aloud data, which provide an "informative window" to expertise in action (deGreef and Breuker 1985). Unfortunately, knowledge acquisition and knowledge base implementation are seldom separated, and thinking aloud data is not widely used in knowledge engineering, primarily because the data is assumed to be difficult to interpret.

3. KNOWLEDGE ACQUISITION: AN EMPIRICAL ASSESSMENT

3.1 Previous Research

There has been very little empirical research conducted in evaluating different knowledge acquisition techniques. Grover (1983) describes an early experiment that evaluated four different knowledge acquisition techniques: forward scenario simulation ("walk throughs"), goal decomposition (20 questions), protocol analysis, and frame analysis. The study found that walk throughs and frame analysis proved most useful. Hoffman (1987) provides anecdotal assessments of different knowledge acquisition techniques. He concludes that all experts, domains, and projects are different, and that some methods will work for some projects and others will not. Prerau (1987) supports this notion and urges knowledge engineers to modify their development strategies to fit the situation and people involved.

3.2 The Prototyping Experiment

The objective of this experiment was to determine what types of heuristic could be elicited through different knowledge acquisition techniques, in order to provide a foundation for the development of a robust, empirically evaluated, and generalizable knowledge acquisition methodology. The experiment provides an empirical assessment of three different knowledge acquisition techniques: scenarios, simulated familiar tasks, and actual familiar tasks. It used as a research vehicle a prototype expert system developed for maritime shipboard piloting applications (Grabowski 1987). Data for the experiment were comprised of expert protocols--from three ship's pilots--gathered as the experts thought aloud, performing duties associated with piloting a ship into and out of New York harbor. Heuristic derived from the experts' protocols were classified and analyzed to determine their contribution to the "expertness" of the expert system, and preliminary guidelines for effecting knowledge acquisition were developed.

3.2.1 The Piloting Expert System

The Piloting Expert System (Grabowski 1987) is a prototype maritime piloting expert system developed for the U.S. Department of Transportation, Maritime Administration, which provides decision support to ship's masters, mates on watch, and pilots navigating in congested harbor waters. The system captures the decision-making expertise of the local pilot and provides local environment (platform, port, weather, visibility, ship-handling, traffic, navigation, etc.) particular ship's piloting recommendations to the operator. In addition, the system allows ship's officers and pilots to mentally rehearse typical and atypical piloting transits before making actual voyages, guided by the collective expertise of the local piloting organization. The system was developed in cooperation with the United New York-New Jersey Sandy Hook Pilot's Organization, the primary ship's piloting organization for New York harbor, and with Puerto Rican Marine Management, Inc., a ship operator whose vessels transit New York harbor on a weekly basis.

The system decreases the information overload under which the ship's pilot presently labors, thus increasing the safety of navigation; provides for more effective distribution of piloting information within the local piloting organization, providing more efficient and consistent knowledge transfer; serves as a training device aboard merchant vessels, pilot boats, and ship simulators, training junior pilots and deck officers; and serves as a voyage planning tool for pilots and deck officers, who now have the luxury of voyage transit rehearsals in realtime.

3.2.2 Experimental Design

The main experiment consisted of collecting protocols of three pilots who thought aloud as they performed duties associated with piloting a ship into and out of New York harbor. The three pilots effected their piloting tasks in three different ways, each way involving a different means of knowledge acquisition.

The first subject was instructed to think aloud while imagining a typical transit (*scenario method*) aboard a 17,500 deadweight ton Roll on/Roll off ship, which was entering and leaving New York harbor. This vessel is typical of those entering the harbor and representative of the type of vessel piloted by New York harbor pilots. The subject was given a current chart of the area, as well as tide tables, navigation tables, and times of sunrise and sunset. Protocols were taken while the subject mentally walked through the transits, accompanied by traditional tools of his trade (charts and nautical publications).

The second subject used the prototype "Piloting Expert System" to simulate a transit of New York harbor with the same 17,500 deadweight ton Roll on/Roll off ship (*simulated familiar task method using a prototype expert system*). This subject was provided with some background familiarization with the expert system, as well as some hands-on training with the system, before being asked to "pilot" the expert system "ship" out of New York harbor. Protocols were taken as the subject thought aloud while using the Piloting Expert System to simulate a transit out of New York harbor.

The last subject was observed as he performed his normal piloting duties aboard the same type of vessel entering New York harbor (*actual familiar task method*). The ship was relatively new (1984), had a full radar and electronic navigational equipment suite, and an electro-hydraulic steering system. All electronic and mechanical navigational and propulsion equipment was fully operational during the transit. The pilot carried with him his radio, his tide book, and a pair of sunglasses. Protocols were again taken as this subject thought aloud as he piloted a vessel into New York harbor.

In each of the three cases, the vessel and transit used were the same. The only differences between the three subjects were the cues provided by the environments in which each effected their task and the means of knowledge acquisition used to elicit their expertise.

3.2.2.1 Subjects

The subjects used for this experiment were experienced (20 years or more) senior Sandy Hook pilots, selected for their participation because of their verbal skills, their piloting expertise, and their interest in the project. All three pilots had joined the Sandy Hook Pilot's Association as pilot apprentices in their teens, been apprentice pilots for the requisite seven years, served as junior pilots for five to ten years (depending on openings available for senior pilot positions), and had been senior pilots for at least seven years. None of the pilots had ever had a ship collision or accident, records that have been achieved over the course of 1,000 transits in New York harbor. Each pilot was recognized by the Sandy Hook Pilot's Association and by his peers as a master professional.

Subject pilots were recommended for participation in the three tasks by the Governing Board of the Sandy Hook Pilot's Association, based on their personal characteristics. The first subject, who participated in the scenario exercise, verbalized thought processes easily and was recommended for the simulated familiar task method because of his verbal and conceptual abilities. The second subject, who participated in the simulated familiar task exercise, was recommended by the Governing Board because of his previous involvement in piloting simulator exercises and his interest in the use of expert systems by piloting organizations. The third subject, who piloted an actual vessel into New York harbor while thinking aloud, was recommended for the actual familiar task role because of his piloting expertise, his ability to verbalize thoughts while engaged in a strenuous activity, and his sense of humor. Only three subjects were utilized for this experiment because of the density of the protocols expected (which require substantial work to transcribe, code, decode, and translate into heuristic) and because of the preliminary nature of the investigation.

3.2.2.2 Procedure

A primary goal of designing this experiment was to maintain a balance between, collecting as much information as possible (which would be useful in interpreting the protocols) and interfering as little as possible with the problem-solving processes reflected in the protocols. The experimental design was formulated after Bouwman (1978), following a review of a variety of protocol analysis experiments, including a previously-collected set of piloting protocols (Huffner 1976). Subjects were first handed a description of the piloting scenario and a set of thinking aloud task instructions. After reviewing the description and the instructions, the subjects then participated in several practice thinking aloud sessions to acquaint them with the process. After the initiation, subjects participated in their respective piloting activities (imagining a transit, using the prototype expert system to simulate a transit, and making an actual piloting transit). Protocols were taken as the subjects thought aloud. Prompts (such as "What are thinking now?" and "Keep talking") were used to prompt verbalizations. Immediately following completion of their tasks, subjects were asked to recall all "significant facts" noticed during the transit, a procedure designed to generate evidence of what information subjects stored in memory during the experiment.

3.2.3 Analysis of the Protocols

The protocol analysis was based upon Waterman and Newell (1971) and Bouwman (1978): episode representations and activity diagrams were used to determine the piloting heuristic. The first step in the process was to transcribe the taped version of the protocol. This audio representation was converted to a lexical representation, which included prosodic features (accents and intonations) and timing information (pauses, protocol density, and syntactic information). These representation were then split into units called topic segments, small pieces of text concerned with only one task topic. Task topics in the piloting domain included particular buoys, a particular ship, a lighthouse, the status of the radar, etc., or single ideas, arguments, or conclusions.

From the topic representation, two new representations were developed: *operators*, which represent objects considered in decision-making, and *relations*, which represent how those objects are manipulated, processed, or thought about. *Episode* representations were developed from a combination of the operators and relations, and heuristic were derived from the operator-relation combinations. Activity diagrams were then constructed from the episode representations. *Activity diagrams* are schematic representations of protocol episodes, with episodes represented as blocks, and the lines between the blocks indicating relationships between episodes. Structuring activity diagrams focused on ascertaining whether adjacent episodes were connected or disconnected and where to attach new episodes in the diagram. Activity diagrams graphically depicted the heuristic derived from the experts.

3.3 Results

Heuristic were divided into two categories: those that were common to all subjects, regardless of knowledge acquisition method, and those that were articulated only by individual subjects, which were considered knowledge acquisition method-specific. In this experiment, 31 percent of all the heuristic were common and 69 percent were knowledge acquisition method-specific. This indicates that, for this experiment, only a third of the piloting expertise was commonly accessible, regardless of the knowledge acquisition method used. This result underscores the importance of utilizing several different means of knowledge acquisition to access the different "packages" of information prompted by different knowledge acquisition techniques (Anderson 1983; Perkins 1981).

Few heuristic (five of a total of 212) were shared between the scenario subject and the actual familiar task subject. In contrast, the scenario subject and the simulated familiar task subject shared 28 common heuristic, while the simulated familiar task and the actual familiar task subjects shared 18. These results are not surprising. One would expect greater correlation between experts dealing with abstract ideas and those using a prototype. Similarly, one would expect less correlation between an expert imagining a transit and one actually performing the task. Not surprisingly, the simulated familiar task expert--

whose task evidenced elements of both abstract and operational natures--shared a number of heuristic with both of the other subjects.

The percentage of method-specific heuristic varied considerably, depending on the leg of the transit. Each subject provided the bulk of his method-specific heuristic at different times in the transit. The scenario expert provided the greatest percentage of his heuristic at the beginning of the experiment, during Legs 1 and 2, while the simulated familiar task subject provided his peak contribution in the middle of the transit. The actual familiar task subject provided his peak percentage of heuristic late in the voyage, when operational concerns and tight turns in the channel became important.

3.3.1 Activity Analysis

The activity analysis also provided interesting results. The *scenario subject* discussed a wide array of potential scenarios, in addition to the specific transit under consideration. He spent a great deal of time conceptualizing ideas, problems, and hypotheses. Because of the nature of the knowledge acquisition method, this subject was not cued to explore or investigate his local environment. The scenario subject focused on operators which subjects in more real-life situations ignored, or paid less attention to. Because of the unstructured manner in which the scenario subject's information was accessed and the operational "slack" the subject experienced, broad *conceptual* domain heuristic were elicited from the scenario subject.

In contrast, the *simulated familiar tasks subject* spent most of his time exploring local cues: checking and exploring local information, summarizing data, and moving on to explore more data. This subject offered few explanations and still fewer judgments, preferring to concentrate instead on the "check, explore, and summarize" pattern. The simulated familiar task subject was also concerned with vessel time, speed, and course, and his heuristic were classified as logistical. The simulated familiar task method allowed the subject to focus on how the decision-making process was effected and whether it was effected in a timely manner. This method allowed the subject to temporarily ignore operational concerns, but still remain linked to the real decision situation. The gap between the operational and conceptual heuristic was bridged by the *logistical* heuristic of the simulated familiar task method.

The verbalizations of the *actual familiar task subject* indicated that he spent most of his time making assessments of where the ship was at any moment and what course he was/should be steering. These heuristic are the essence of operational piloting. This subject had little time for checking, exploring, or problem-hypothesizing. His task was to pilot the vessel safely and in a timely fashion. As a result, the actual familiar task subject focused on the ship's position and its course, with little attention devoted to the ship's handling characteristics, ship's equipment status, or navigational aids. The heuristic gleaned from the actual familiar task subject's protocols were of an *operational* nature. The subject concentrated on the task at hand and the associated heuristic reflected this orientation.

4. CONCLUDING COMMENTS

There are several implications of the research. First, in this experiment, 30 percent of the experts' heuristic were common, indicating that a set of piloting heuristic was accessible, regardless of the knowledge acquisition method employed. More importantly, the experiment demonstrates that a majority (in this experiment, 70 percent) of the expert heuristic were *not* common and, thus, were knowledge acquisition method-dependent. This result provides empirical support for the anecdotal recommendation to use several knowledge acquisition techniques in order to access different "packages" of knowledge.

The research also provides an initial framework that categorizes the types of information developers might expect to gather when using different means of knowledge acquisition. When experts use scenarios as a knowledge acquisition technique, *conceptual* heuristic can be expected; *operational* heuristic can be expected to be elicited by actual familiar task performance; and simulated familiar task methods can be expected to yield *logistical* heuristic.

This research is a first step in a broader knowledge acquisition research program. In this set of experiments, the types of heuristic which can be expected from three different knowledge acquisition techniques were investigated. A number of questions fall from this small, preliminary experiment:

- What kinds of heuristic can be elicited from *each* of the different knowledge acquisition techniques discussed--not just constrained processing tasks?
- How generalizable are the results of this preliminary study? To what extent is this a useful piloting study, rather than the foundation of generalized knowledge acquisition experimental design?
- How do differences in experts affect the generalizability of the experiment's results? To what extent do individual differences play a role in empirical knowledge acquisition results?
- How does the experiment's small sample size affect the findings?
- How do the findings in this study, in the maritime piloting domain, apply to other situation assessment domains? Are the results applicable outside the

realm of first generation situation assessment expert systems?

The next step in the broader knowledge acquisition research program is to investigate the questions raised by the limited preliminary study, and to use the results as the foundation for a robust, pragmatic knowledge acquisition methodology. The goal of this research program is the development of a general knowledge acquisition methodology which can be experimentally evaluated. The results of the evaluation will provide guidance to researchers and practitioners in selecting knowledge acquisition methods (or combinations of methods), based on the type of problem being addressed, the number(s) and type(s) of experts required, the depth of knowledge of the problem domain, and the resources available for development of the system.

5. ACKNOWLEDGEMENTS

This research was supported in part by the United States Department of Transportation, Maritime Administration.

Martha Grabowski is Visiting Assistant Professor, Decision Sciences and Engineering Systems, Rensselaer Polytechnic Institute, Troy, New York, and Assistant Professor, Management Information Systems, LeMoyne College, Syracuse, New York.

6. REFERENCES

Alavi, M. "An Assessment of the Prototyping Approach to Information Systems Development." *Communications of the ACM*, 27:6, June 1984, pp. 556-563.

Anderson, J. R. *The Architecture of Cognition*. Cambridge, Massachusetts: Harvard University Press, 1983.

Belkin, N. J.; Brooks, H. M.; and Daniels, P. J. "Knowledge Elicitation and Discourse Analysis." *Proceedings of the Knowledge Acquisition for Knowledge-Based Systems Workshop*, Banff, Alberta, Canada, 2-7 November 1986, pp. 3-0-3-16.

Bimson, K. D., and Burris, L. B. "The Craft of Engineering Knowledge: Some Practical Insights." *Twentieth Annual Hawaii International Conference on System Sciences*, 1987, pp. 460-469.

Bouwman, M. "Financial Diagnosis:A Cognitive Model of the Processes Involved." Unpublished Ph.D. Thesis, Carnegie-Mellon University, 1978.

Brooking, A. G. "The Analysis Phase in the Development of Knowledge-Based Systems." In W. A. Gale (ed.), *Artificial Intelligence and Statistics*, Reading, MA: Addison-Wesley Publishing Company, 1986, pp. 321-334.

Buchanan, B. G.; Barstow, D.; Bechtel, R.; Bennett, J.; Clancey, W.; Kulikowski, C.; Mitchell, T. M.; and Waterman, D. A. "Constructing Expert Systems." In F. Hayes-Roth, D. A. Waterman, and D. B. Lenat (eds.), *Building Expert Systems*. Reading, MA: Addison-Wesley Publishing Co., 1983, pp. 127-168.

Buchanan, B. G., and Shortliffe, E. H. *Rule-Based Expert Systems*. Reading, MA: Addison-Wesley Publishing Company, 1984.

Butler, K. A., and Corter, J. E. "Use of Psychometric Tools for Knowledge Acquisition: A Case Study." In W. A. Gale (ed.), *Artificial Intelligence in Statistics*. Reading, MA: Addison-Wesley Publishing Co., 1986, pp. 295-319.

Cooke, N. M., and McDonald, J. E. "A Formal Methodology for Acquiring and Representing Expert Knowledge." *Proceedings of the IEEE*, 74:10, October 1986, pp. 1422-1430.

Davis, G. B., and Olson, M. H. *Management Information Systems*, Second Edition. New York: McGraw-Hill Book Co., 1985.

Davis, R.; Austin, H.; Carlbom, Frawley, B.; Pruchnik, P.; Sneiderman, R.; and Gilreath, J. "The DIPMETER ADVISER: Interpretation of Geologic Signals." *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, 1981, pp. 846-849.

Davis, R.; Buchanan, B. G.; and Shortliffe, E. H. "Production Rules as a Representation for a Knowledge-Based Consultation Program." *Artificial Intelligence*, 8 (1977), pp. 15-45.

deGreef, P., and Breuker, J. "A Case Study in Structured Knowledge Acquisition." *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, 1985, pp. 390-392.

Eliot, L. B. "Knowledge Acquisition from Multiple Experts for Developing Expert Systems and Decision Support Systems: Use of Multidimensional Scaling." University of Southern California Working Paper, 1986.

Ericsson, K. A., and Simon, H. A. "Verbal Reports as Data." *Psychological Review*, No. 87, May 1980, pp. 215-251.

Fellers, J. W. "Skills and Techniques for Knowledge Acquisition: Survey, Assessment, and Future Directions." In J. I. DeGross and C. H. Kriebel (eds.), *Proceedings of the Eighth International Conference on Information Systems*, Pittsburgh, December 1987, pp. 118-132.

Gammack, J. G., and Young, R. M. "Psychological Techniques for Eliciting Expert Judgment." In M. A. Bramer (ed.), Research and Development in Expert Systems, London: Cambridge University Press, 1985, pp. 105-112.

Grabowski, M. R. "Prototyping As a Design Strategy and a Means of Knowledge Acquisition for Expert Systems: An Assessment Using 'The Piloting Expert System'." Unpublished Ph.D. Thesis, Rensselaer Polytechnic Institute, Troy, New York 1987.

Grabowski, M. R., and Wallace, W. A. "Prototyping Expert Systems." School of Management Research paper 37-P6-86, Rensselaer Polytechnic Institute, Troy, New York, 1986.

Greene, D. P. "Automated Knowledge Acquisition: Overcoming the Expert Systems Bottleneck." In *J. I. DeGross and C. H. Kriebel (eds.), Proceedings of the Eighth International Conference on Information Systems*, Pittsburgh, December 1987, pp. 107-117.

Grover, M. D. "A Pragmatic Knowledge Acquisition Methodology." *Proceedings of the Eighth International Joint Conference on Artificial Intelligence* (IJCAI-83), Karlsruhe, West Germany, 1983, pp. 436-438.

Hayes-Roth, F.; Waterman, D. A.; and Lenat, D. B. "An Overview of Expert Systems." In F. Hayes-Roth, D. A. Waterman and D. B. Lenat (eds.), *Building Expert Systems*. Reading, MA: Addison-Wesley Publishing Company, 1983, pp. 3-29.

Hoffman, R. H. "The Problems of Extracting the Knowledge of Experts from the Perspective of Experimental Psychology." *AI Magazine*, 8:2, Summer 1987, pp. 53-67.

Huffner, J. R. "Pilotage in Confined Waterways of the U.S.: A Preliminary Study of Piloting Decision-Making." U.S. Department of Transportation, US Coast Guard, Office of Research and Development, Report No. CG-D-96-76, July 1976.

Janson, M. A., and Smith, L. D. "Prototyping for Systems Development: A Critical Appraisal." *MIS Quarterly*, December 1985, pp. 305-316.

Newell, A., and Simon, H. A. *Human Problem Solving*, Englewood Cliffs, NJ: Prentice-Hall, Inc., 1972.

Perkins, D. N. *The Mind's Best Work*. Cambridge, MA: Harvard University Press, 1981.

Prerau, D. S. "Selection of an Appropriate Domain for an Expert System." *AI Magazine*, 6:2, Summer 1985, pp. 26-30.

Raulefs, P. "Knowledge Processing for Expert Systems." In T. Bernald and G. Albers (eds.), *Artificial Intelligence: Towards Practical Application*. Amsterdam: North Holland Publishing Company, 1985.

Rolandi, W. G. "Knowledge Engineering in Practice." *AI Expert*, 1:4, December 1986, pp. 58-62

Schweiger, D. M. "Is the Simultaneous Verbal Protocol a Viable Method for Studying Managerial Problem-Solving and Decision-Making?" *Academy of Management Journal*, 26:1, August 1983, pp. 185-192.

Vesonder, G. T.; Solfo, S. J.; Zielinski, J. E.; Miller, F. D.; and Copp, D. H. "ACE: An Expert System for Telephone Cable Maintenance." *The Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, 1983, pp. 116-123.

Waterman, D. A. *A Guide to Expert Systems*. Reading, MA: Addison-Wesley Publishing Company, 1986.

Waterman, D. A., and Newell, A. "Protocol Analysis as a Task for Artificial Intelligence." *Artificial Intelligence*, No. 2, 1971, pp. 285-318.

Whalley, P. C. "The Psychology of Similarity." Unpublished Ph.D. Thesis, Open University, 1984.

Wright, G., and Ayton, P. "Eliciting and Modeling Expert Knowledge." *Decision Support Systems*, 3 (1987), pp. 13-26.